

WHAT IS CLAIMED IS:

1. A storage system comprising:
 - 5 a storage array controller; and
 - a storage array coupled to the storage array controller and comprising:
 - a plurality of disk drives; and
 - a plurality of disk drive controllers, wherein each disk drive controller is
 - 10 configured to control one or more of the disk drives;
 - wherein the storage array controller is configured to control accesses to data within the storage array;
 - 15 wherein the storage array controller is further configured to issue a first scrubbing operation command to a first disk drive controller;
 - wherein the first disk drive controller is configured to receive the first scrubbing operation command from the storage array controller and, in response to
 - 20 receiving the first scrubbing operation command, to read data from within a data range from at least one of the disk drives, to calculate a new checksum for the data, and to compare the new checksum to a preexisting checksum for the data; and
 - 25 wherein if the new checksum differs from the preexisting checksum, the first disk drive controller is further configured to respond to the first scrubbing operation command by indicating that the data within the data range is erroneous.

2. The storage system of claim 1, wherein the storage array controller is configured to specify the data range in the first scrubbing operation command.
3. The storage system of claim 2, wherein the storage array controller is configured to specify the data range by specifying a starting address of the data range and a number of blocks within the data range.
4. The storage system of claim 1, wherein the storage array controller is further configured to initiate a reconstruction attempt of the erroneous data in response to the first disk drive controller indicating that the data is erroneous.
5. The storage system of claim 1, wherein the first disk drive controller is further configured to respond to the first scrubbing operation command by indicating an address of the erroneous data.
6. The storage system of claim 1, wherein the storage array controller is further configured to calculate the preexisting checksum for the data when the data is written to the storage array.
7. The storage system of claim 1, wherein the storage array controller is configured to specify the preexisting checksum in the first scrubbing operation command.
8. The storage system of claim 1, wherein the storage array controller is further configured to receive the preexisting checksum for the data from a host computer system when the host computer system writes the data to the storage array.
9. The storage system of claim 1, wherein the storage array controller is further configured to maintain the preexisting checksum in a cache coupled to the storage array controller.

10. The storage system of claim 1, wherein the first disk drive controller comprises a checksum engine configured to calculate the new checksum.
11. The storage system of claim 1, wherein the first disk drive controller is further
5 configured to indicate that a scrubbing operation was successful if the new checksum equals the preexisting checksum.
12. The storage system of claim 1, wherein the storage array controller is further
10 configured to issue multiple scrubbing operation commands in parallel to multiple ones of the disk drive controllers.
13. The storage system of claim 1, wherein the storage array controller is implemented in software running on a host computer system.
14. The storage system of claim 1, wherein the first disk drive controller comprises at
15 least one second disk drive controller integrated with the at least one of the disk drives and a third disk drive controller coupled between the storage array controller and the at least one second disk drive controller, wherein the at least one second disk drive controller is configured to read the data within the data range from the at least one of the
20 disk drives and to provide the data to the third disk drive controller, and wherein the third disk drive controller is configured to calculate the new checksum.
15. The storage system of claim 1, wherein the storage array controller is further
25 configured to store the preexisting checksum in a checksum location that is independent of a location where the data within the data range is stored.
16. The storage system of claim 1, wherein the data range is a multi-unit data range, and wherein the storage array controller is further configured to calculate the preexisting checksum directly from the data in the multi-unit data range.
- 30

17. The storage system of claim 1, wherein the data range is a multi-unit data range, and wherein the storage array controller is further configured to calculate the preexisting checksum as an extent checksum computed from a plurality of unit checksums that each correspond to a unit within the multi-unit data range.

5

18. The storage system of claim 17, wherein the multi-unit data range is a multi-byte data range.

10

19. The storage system of claim 17, wherein the multi-unit data range is a multi-block data range.

20. The storage system of claim 17, wherein the storage array controller is further configured to calculate the preexisting checksum from the plurality of first unit checksums when issuing each scrubbing operation command.

15

21. The storage system of claim 1, wherein the storage array controller is further configured to issue a second scrubbing operation command to the first disk drive controller in response to the first disk drive controller indicating that the data within the data range is erroneous, and wherein, in response to receiving the second scrubbing operation command, the first disk drive controller is configured to read second data within a second data range, calculate a new second checksum for the second data within the second data range, and to compare the new second checksum to a preexisting second checksum, wherein the second data range is a subset of the first data range.

20

22. The storage system of claim 21, wherein the storage array controller is further configured to specify the second data range and the second preexisting checksum in the second scrubbing operation command.

25

23. A method of performing scrubbing operations in a storage array, comprising:

30

a storage array controller issuing a first scrubbing operation command to a disk controller;

5 in response to receiving the first scrubbing operation command, the disk controller reading the data within a first data range, computing a new first checksum for the data within the first data range, and comparing the new first checksum to an original first checksum for the first data range; and

10 if the new first checksum differs from the original first checksum, the disk controller indicating that the data within the first range is erroneous.

24. The method of claim 23, wherein the first scrubbing operation command specifies the first data range.

15 25. The method of claim 24, wherein the first scrubbing operation command specifies the first data range by specifying a starting address of the first data range and a number of blocks within the first data range.

20 26. The method of claim 23, wherein said responding further comprises indicating an address of the erroneous data.

27. The method of claim 23, further comprising initiating a reconstruction attempt of the erroneous data in response to said indicating.

25 28. The method of claim 23, further comprising the storage array controller issuing a second scrubbing operation command to the disk controller in response to said indicating.

29. The method of claim 28, wherein the second command specifies a second data range and an original second checksum for data within the second data range, and
30 wherein the second data range is a subset of the first data range.

30. The method of claim 29, further comprising the disk controller computing a new second checksum for the second data range, comparing the new second checksum to the original second checksum, and indicating that the data within the second data range is erroneous if the new second checksum does not equal the original second checksum.

31. The method of claim 28, further comprising calculating the original first checksum using a first checksum algorithm and calculating the original second checksum using a second checksum algorithm, wherein the first checksum algorithm is a stronger checksum algorithm than the second checksum algorithm.

32. The method of claim 28, further comprising the storage array controller issuing a plurality of additional scrubbing operation commands to the disk controller in response to said indicating, wherein the first data range is a multi-unit data range, and wherein each additional scrubbing operation command specifies a unit of data within the first data range.

33. The method of claim 23, further comprising the disk controller indicating that the scrubbing operation was successful if the new first checksum equals the original first checksum.

34. The method of claim 23, further comprising the storage array controller issuing a plurality of first scrubbing operation commands in parallel to a plurality of disk controllers.

35. The method of claim 23, wherein said issuing further comprises the storage array controller providing the original first checksum to the disk controller.

36. The method of claim 23, further comprising:

a host computer system writing the data within the first data range to the storage array, wherein said writing comprises the host computer system providing the original first checksum; and

5 storing the original first checksum within the storage array in response to said writing.

37. The method of claim 23, further comprising the storage array controller calculating the original first checksum in response to the data within the first data range
10 being written to the storage array by a host computer system.

38. The method of claim 37, wherein said calculating the original first checksum comprises calculating the original first checksum from the data within the first data range.

15 39. The method of claim 23, wherein the storage array controller is implemented in software running on a host computer system.

40. The method of claim 23, wherein the disk controller comprises at least one first disk controller integrated with a disk drive and a second disk controller coupled between
20 the storage array controller and the at least one first disk controller, wherein the at least one first disk controller is configured to read the data within the first data range from the disk drive and to provide the data to the second disk controller, and wherein the second disk controller is configured to calculate the new first checksum.

25 41. The method of claim 23, further comprising calculating the original first checksum and storing the original first checksum in a checksum location that is independent of a location where the first data within the first data range is stored.

42. The method of claim 23, wherein the first data range is a multi-unit data range, and wherein the original first checksum comprises a first extent checksum computed directly from data in the multi-unit data range.
- 5 43. The method of claim 23, wherein the first data range is a multi-unit data range, and wherein the original first checksum comprises a first extent checksum computed from a plurality of first unit checksums that each correspond to a unit within the multi-unit data range.
- 10 44. The method of claim 43, wherein the multi-unit data range is a multi-byte data range.
45. The method of claim 43, wherein the multi-unit data range is a multi-block data range.
- 15 46. The method of claim 43, further comprising the array controller calculating the first extent checksum from the plurality of first unit checksums on the fly when issuing the first scrubbing operation command.
- 20 47. A disk controller, comprising:
- a first interface configured to be connected to a storage array via an interconnect;
- a second interface configured to access a storage media; and
- 25 a controller configured:
- to receive a scrubbing operation command via the first interface;

in response to receiving the scrubbing operation command, to read data within a range of data from the storage media;

to calculate a checksum for the data;

5

to compare the checksum to a preexisting checksum for the data; and

if the checksum differs from the preexisting checksum, to indicate that the data within the data range is erroneous.

10

48. The disk controller of claim 47, wherein the controller is configured to access the storage media by communicating the data range to an integrated disk controller coupled to the storage media and by receiving the data from the integrated disk controller.

15 49. The disk controller of claim 47, wherein the scrubbing operation command specifies the data range.

50. The disk controller of claim 49, wherein the scrubbing operation command specifies the data range by specifying a starting address of the data range and a number of
20 blocks within the data range.

51. The disk controller of claim 47, wherein the controller is further configured to indicate an address of the erroneous data if the checksum differs from the preexisting checksum.

25

52. The disk controller of claim 47, wherein the scrubbing operation command specifies the preexisting checksum.

53. The disk controller of claim 47, further comprising a checksum engine configured
30 to calculate the checksum.

54. The disk controller of claim 47, wherein the controller is further configured to indicate that the scrubbing operation was successful if the checksum equals the preexisting checksum.

5